

National Strategic Computing Initiative Strategic Plan

July 2016

Report prepared by

THE NATIONAL STRATEGIC COMPUTING INITIATIVE EXECUTIVE COUNCIL

Co-Chair

John P. Holdren

Assistant to the President for Science
and Technology and Director,
Office of Science and Technology Policy

Co-Chair

Shaun Donovan

Director,
Office of Management and Budget

Members

Department of Defense
Department of Energy
Department of Homeland Security
Federal Bureau of Investigation
Intelligence Advanced Research Projects Activity
National Aeronautics and Space Administration
National Institutes of Health
National Institute of Standards and Technology
National Oceanic and Atmospheric Administration
National Science Foundation
National Security Agency

Additional Participants

National Economic Council
National Security Council
Office of the Director for National Intelligence

Table of Contents

Table of Contents

Executive Summary.....	3
1. Introduction and Overview	5
2. Roles and Responsibilities.....	6
2.1. Agency Roles	7
2.2. Opportunities for Academic and Industry Engagement	8
3. Achieving the NSCI Strategic Objectives	10
3.1. Strategic Objective 1: Capable Exascale	10
3.2. Strategic Objective 2: Technology Coherence	13
3.3. Strategic Objective 3: Computing Beyond Moore’s Law	14
3.4. Strategic Objective 4: An Enduring National HPC Ecosystem	16
3.5. Strategic Objective 5: Advancement through Public-Private Collaboration.....	19
4. Summary	21
A. Annex A: Related Initiatives	22
A.1 Advanced Manufacturing Initiatives.....	22
A.2 BRAIN Initiative.....	22
A.3 Computer Science for All	22
A.4 Materials Genome Initiative	23
A.5 National Big Data R&D Initiative.....	23
A.6 National Nanotechnology Initiative.....	23
A.7 Precision Medicine Initiative	24
B. Annex B: E.O. 13702.....	25

Executive Summary

The National Strategic Computing Initiative (NSCI) is a whole-of-Nation effort to sustain and enhance U.S. leadership in high-performance computing (HPC). The NSCI seeks to accomplish five strategic objectives in a government collaboration with industry and academia: (1) accelerate the successful deployment and application of capable exascale computing; (2) ensure that new technologies support coherence in data analytics as well as simulation and modeling; (3) explore and accelerate new paths for future computing architectures and technologies, including digital computing and alternative computing paradigms; (4) holistically expand capabilities and capacity of a robust and enduring HPC ecosystem; and (5) establish an enduring public-private collaboration to ensure shared benefit across government, academia, and industry.

Objective 1: The NSCI seeks to accelerate the development of HPC systems with the goal of achieving capable exascale computing by the mid-2020s to meet the needs of critical government applications and capitalize on emerging datasets. Although custom HPC systems that achieve peak performance of an exaflop on specialized benchmarks could be realized by the end of this decade, these systems would not be able to sustain performance at that level on realistic calculations, be expensive to power, be very difficult to program, lack the memory capacity to perform current and anticipated new classes of scientific and engineering applications, and be potentially unreliable in performing long calculations. Exascale computing systems with important characteristics such as affordable power consumption, programmability, reliability, and adequate memory and networking are not projected to become available commercially until the mid to late 2020s, according to industry roadmaps.

Objective 2: The NSCI seeks to develop a coherent platform for modeling, simulation, and data analytics, primarily through the development of a more agile and reusable HPC software portfolio. Historically, there has been a separation between data analytic computing and modeling and simulation. Systems have been optimized for a specific class of applications, but the differences between these application spaces are fading rapidly. The growth of extremely large-scale data analytics within the modeling and simulation community demands a dynamic interaction between analysis and simulations. As data analytics increases in computational intensity, and modeling and simulation encounter increasingly complex problems, both fields face barriers to scalability along with new demands for interoperability, robustness, and reliability of results.

Objective 3: To develop new computing capabilities, the NSCI will pursue two parallel lines of effort over a 10-20 year period: (1) the research and development (R&D) of technologies that will move digital computing performance past the theoretical limits of complementary metal-oxide semiconductors, and (2) the R&D of alternative paradigms that will open up new possibilities for large-scale computing. While current semiconductor approaches should be feasible at exascale, these approaches will eventually plateau due to the physical limitations inherent in semiconductor technologies. It would be premature to select individual technology paths for alternative computing paradigms or alternative digital computing technologies, so the NSCI will initially focus on broad research to identify areas of promise across the span of future computing technologies.

Objective 4: The NSCI seeks to develop and adopt new approaches, technologies, and software architectures to support HPC application development, reusability, trustworthiness, sustainability, and HPC workforce development. NSCI also seeks to increase access to HPC technologies and platforms. The current HPC ecosystem of software, hardware, networks, and workforce is neither widely available nor sufficiently flexible to support emerging opportunities. The need for change is driven by the growing complexity and size of a broad array of simulations; the expansion of the role of large-scale computation

in emerging frontiers of science; and the need for dynamic interaction of computation with other elements of the cyberinfrastructure, including distributed participants, instruments, large data repositories, high-bandwidth networks, and mobile devices. Greater accessibility and interweaving with production environments, however, will increase exposure of HPC capabilities to the full spectrum of cyber threats.

Objective 5: The NSCI seeks to broaden public-private collaboration for advancing HPC technologies to support Federal missions, scientific discovery, and economic competitiveness. From the vector-based supercomputer era to the massively parallel systems of modern HPC, collaboration between Federal agencies and the private sector drove the transformations that define today's international HPC landscape. The inclusive nature of this public-private collaboration endeavors to benefit the public and private sectors, across missions for scientific discovery, economic competitiveness, and national security.

Realizing the vision of the NSCI will demand a fully developed HPC ecosystem that meets the needs of government, industry, and academia. Some aspects of that ecosystem have broad commercial drivers and will naturally evolve from the current semiconductor environment. This Strategic Plan (Plan) focuses on areas where government engagement is essential in creating the technological capability, computational foundations, and workforce capacity to realize the vision of the NSCI. The Plan identifies the roles assigned to Federal agencies, and highlights ongoing and planned activities that will contribute to NSCI's goals.

The combination of broad commercial drivers and government action will not be sufficient to achieve the vision of the NSCI. The success of the NSCI depends upon development, commercialization, and deployment of new HPC technologies and infrastructure. The NSCI strives to establish collaboration among the Federal Government, industry, and academia to develop the technological capabilities and computational foundations that will support scientific discovery and economic drivers for the 21st century, and that will not naturally evolve from current commercial activity.

1. Introduction and Overview

HPC is essential to the Nation's global economic competitiveness and scientific discovery. Over the past 60 years, U.S. leadership in computing has been maintained through continual development and deployment of new computing technology with increased performance on applications of real-world significance. Sustaining this leadership in HPC in the coming decades requires a national response to increasing computing demands, emerging technological challenges, and growing international competition. To meet these challenges and maximize the benefits of HPC, the President established the National Strategic Computing Initiative (NSCI), an effort to create a cohesive, multi-agency strategic vision and Federal investment strategy in HPC in collaboration with industry and academia.

This whole-of-Nation approach reflects the critically important role that academia and industry will play in the NSCI. The guiding principles and strategic objectives of the NSCI, as well as the roles and responsibilities of the NSCI agencies, were specified in Executive Order (E.O.) 13702: *Creating a National Strategic Computing Initiative*.¹

NSCI is guided by four overarching principles. As befitting a whole-of-Nation effort, these principles are stated in E.O. 13702 as imperatives for the entire nation. The principles are presented below:

1. The United States must deploy and apply new HPC technologies broadly for economic competitiveness and scientific discovery.
2. The United States must foster public-private collaboration, relying on the respective strengths of government, industry, and academia to maximize the benefits of HPC.
3. The United States must adopt a whole-of-government approach that draws upon the strengths of and seeks cooperation among all executive departments and agencies with significant expertise or equities in HPC while also collaborating with industry and academia.
4. The United States must develop a comprehensive technical and scientific approach to transition HPC research on hardware, system software, development tools, and applications efficiently into development and, ultimately, operations.

E.O. 13702 identified five strategic objectives that are essential to sustaining and enhancing U.S. scientific, technological, and economic leadership in HPC. The objectives are presented below in an abbreviated form²:

1. Accelerate delivery of a capable exascale computing system delivering approximately 100 times the performance of current systems across a range of applications.
2. Increase coherence between the technology base used for modeling and simulation and that used for data analytic computing.
3. Establish a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post-Moore's Law era").
4. Increase the capacity and capability of an enduring national HPC ecosystem.

¹ *Executive Order Creating a National Strategic Computing Initiative*, E.O. 13702, July 29, 2015, <https://www.whitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-initiative>, July 29, 2015.

² The full text of each objective may be found in Section 3 of this Plan, as well as in E.O. 13702.

5. Develop an enduring public-private collaboration to ensure that the benefits of the research and development advances are shared among government, industrial, and academic sectors.

This Plan describes how the NSCI agencies intend to fulfill the guiding principles, accomplish the strategic objectives, and sustain U.S. leadership in HPC.³ The Plan provides a mechanism to identify, support, and align efforts of NSCI agencies with efforts in academia and the private sector.

During Fiscal Year (FY) 2016, agencies are implementing the strategy set forth in E.O. 13702 through previously planned HPC research, development, and deployment activities. Agencies are also preparing for future activities that will support the principles and objectives of the NSCI. The President's Budget for FY 2017 supports NSCI investments across the Federal Government, including key investments within the Department of Energy (DOE) (\$285 million) and the National Science Foundation (NSF) (\$33 million).

U.S. leadership in HPC is not an end in itself—it provides a powerful capability to help achieve national priorities such as global economic competitiveness, scientific discovery, and national security. To maximize the impact of NSCI, it will have a synergistic relationship with many other national initiatives, integrating scientific findings and technologies, providing HPC and large-scale analytic computing, co-developing advanced software applications, and training students and the workforce to take full advantage of HPC's expanded capabilities. Examples of these national initiatives include:

- Advanced Manufacturing Initiative: <https://www.manufacturing.gov>
- BRAIN Initiative: <https://www.whitehouse.gov/BRAIN>
- Computer Science for All: <https://www.whitehouse.gov/blog/2016/01/30/computer-science-all>
- Materials Genome Initiative: <https://www.mgi.gov>
- National Big Data R&D Initiative: [https://www.nitrd.gov/nitrdgroups/index.php?title=Big_Data_\(BD_SSG\)](https://www.nitrd.gov/nitrdgroups/index.php?title=Big_Data_(BD_SSG))
- National Nanotechnology Initiative: <http://www.nano.gov>
- Precision Medicine Initiative: <https://www.whitehouse.gov/precision-medicine>

The Plan reflects the basic strategy outlined in E.O. 13702, with a particular focus on the NSCI strategic objectives identified therein. [Section 2](#) of this Plan describes NSCI governance and the roles and responsibilities of agencies involved in the NSCI. [Section 3](#) provides details of the supporting whole-of-nation actions and agency contributions needed to address the strategic objectives. [Section 4](#) provides a summary and conclusions.

2. Roles and Responsibilities

To achieve the goals of the NSCI, E.O. 13702 established an Executive Council (EC) to ensure accountability for and coordination of related research, development, and deployment. The EC is co-chaired by the Director of the Office of Science and Technology Policy (OSTP) and the Director of the Office of Management and Budget (OMB), includes agency principals, and meets no less than twice yearly. The EC authored this Plan to increase awareness of government efforts towards the goals of the NSCI, and to maximize public-private engagement with academia and industry.

³ In this document, references to planned programs are offered solely to enhance coordination with academia and the private sector, as directed in EO 13702. These references should not be construed to imply an actual funding request or a commitment from any individual agency to make such a request.

2.1. Agency Roles

E.O. 13702 identifies three broad categories (see Table 1) of agencies currently engaged in HPC research, development, and deployment: lead agencies, foundational R&D agencies, and deployment agencies. These groups may evolve over time as HPC-related mission needs emerge.

Lead agencies include the Department of Defense (DOD), DOE, and NSF, which will engage in mutually supportive R&D within their respective missions in hardware and software, as well as in developing the workforce to support the objectives of the NSCI and to develop and deliver the next generation of integrated HPC capabilities. DOD will focus on data analytic computing to support its missions. Within DOE, the Office of Science and the National Nuclear Security Administration (NNSA) will execute a joint program to develop and deploy advanced simulation through a capable exascale computing program to support their missions. NSF will play a central role in scientific discovery, the broader HPC ecosystem for scientific discovery, and workforce development, which cuts across the NSCI objectives.

Agency	Lead Agencies	Foundational R&D Agencies	Deployment Agencies
Department of Defense (DOD)	X		
Department of Energy (DOE)	X		
Department of Homeland Security (DHS)			X
Federal Bureau of Investigation (FBI)			X
ODNI/Intelligence Advanced Research Projects Activity (IARPA)		X	
National Aeronautics and Space Administration (NASA)			X
National Institutes of Health (NIH)			X
National Institute of Standards and Technology (NIST)		X	
National Oceanic and Atmospheric Administration (NOAA)			X
National Science Foundation (NSF)	X		

Table 1. Summary of Agency Roles

Foundational R&D agencies include the Intelligence Advanced Research Projects Activity (IARPA) and the National Institute of Standards and Technology (NIST), which are charged with fundamental scientific discovery work and associated advances in engineering necessary to support the NSCI objectives. IARPA will focus on future computing paradigms offering an alternative to standard semiconductor computing technologies. NIST will focus on measurement science to support future computing technologies. The foundational R&D agencies will coordinate with academia and industry in conducting this R&D, and with industry and deployment agencies to enable effective transition of R&D efforts into commercial offerings that support the wide variety of requirements across the Federal Government.

Deployment agencies will develop mission-based HPC requirements to influence the early stages of the design of new HPC systems, software, and applications; integrate the special requirements of their respective missions; and engage with the private sector and academia to develop corresponding target HPC requirements through bi-lateral engagements and Requests for Information (RFIs). There are six deployment

agencies: the National Aeronautics and Space Administration (NASA), the Federal Bureau of Investigation (FBI), the National Institutes of Health (NIH), the Department of Homeland Security (DHS), and the National Oceanic and Atmospheric Administration (NOAA). Deployment agencies will also participate in testing, supporting workforce development activities, and engaging with industry to ensure effective deployment within agency mission contexts.

The lead agencies will work closely with the foundational R&D agencies to develop and deploy “post-Moore’s Law” computing technologies. The lead agencies will also work with the foundational R&D agencies and the deployment agencies to increase technology coherence, promote collaboration with industry and academia, and to ensure the NSCI addresses agency mission requirements, while considering a wide variety of national needs. Through coordination with OSTP and OMB, all agencies participating in the NSCI will ensure that NSCI activities align with, leverage, or otherwise engage in mutually supportive efforts with other national strategic initiatives wherever possible. Selected initiatives and their relationships with the NSCI are described in Annex A.

2.2. Opportunities for Academic and Industry Engagement

E.O. 13702 assigns roles and responsibilities solely to government agencies, but recognizes the importance of public-private collaboration to achieving the full potential of NSCI.

The NSCI envisions:

- New hardware and software technologies to support the next generation of digital HPC applications;
- New hardware and software technologies to support the new computing paradigms;
- A robust national HPC ecosystem that supports emerging classes of commercial and scientific applications;
- A larger and more skilled HPC workforce that can take advantage of emerging technologies, including capabilities to support massive-concurrency, data-intensive workflows, tightly-coupled applications, and time-critical responses;
- Economic benefits for companies—large, medium, and small—that integrate use of HPC tools into their product development and manufacturing cycles; and
- National benefits through enhanced scientific discovery across the U.S. research enterprise in universities, industries, government, and national laboratories.

Realizing the vision of the NSCI will require a diverse, dynamic, and highly interoperable HPC ecosystem. The NSCI focuses on those areas where government engagement with industry and academia will spur the development of technological capabilities, computational foundations, and workforce capacity. Broadly available HPC resources will require innovative approaches for development of an HPC ecosystem, including public open-source and privately owned and operated applications developed by either public or private sectors.

Examples of mechanisms available to private-sector organizations to participate in the NSCI include:

- Participating in NSCI and NSCI-related workshops;
- Responding to Requests For Information (RFIs) on NSCI or NSCI-related technologies;
- Reviewing NSCI publications—see *Networking and Information Technology Research and Development (NITRD) website for NSCI*⁴;

⁴ <https://www.nitrd.gov/nsci>

- Establishing or expanding research investments in NSCI-related topics;
- Submitting innovative research proposals in response to grants solicitations, Small Business Innovation Research (SBIR) programs, etc.;
- Partnering with agencies to develop and fund joint solicitations;
- Participating in government-sponsored co-design efforts;
- Participating in joint R&D efforts with universities;
- Expanding classroom and online course offerings in computational science and engineering; and
- Supporting workforce-development activities that directly or indirectly support the NSCI.

Upcoming opportunities for engagement to advance the goals of the NSCI will be available at the NITRD website.⁵

3. Achieving the NSCI Strategic Objectives

This section describes the planned approaches for achieving the five NSCI strategic objectives defined in E.O. 13702, including ties and interdependencies among the objectives, the agencies' respective roles and shared activities for the next 2 years, and opportunities for participation and collaborations by the private sector.

3.1. Strategic Objective 1: Capable Exascale

Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.

Box 1: Capable Exascale Computing

Over the past decade, the United States has developed and deployed computer systems that provide sustained petascale (10^{15} operations per second) performance in scientific and engineering simulations. Petascale-class computer systems are operating today across DOD, DOE, NOAA, NASA, and several NSF-supported supercomputer facilities, providing essential capabilities for science and national security missions. Petascale-class computers also play a key role in U.S. economic competitiveness, with applications in many commercial sectors, including aerospace, pharmaceuticals, and petroleum. The next major performance target, exascale computing (10^{18} operations per second), could be realized in a few years with extensions of the existing technology base in microelectronics. Such a system, however, would be expensive to operate, difficult to program, and would lack the networking and storage capacity to perform new classes of applications.

The NSCI sets a more ambitious goal, achieving *capable exascale computing* to ensure that the performance increase extends beyond benchmarks. Capable exascale computing will require scaling algorithms to billion-way parallelism, improving energy efficiency, and speeding memory access. According to industry roadmaps, exascale computing that addresses these significant challenges will become available through normal commercial drivers by 2030. To support critical government applications and capitalize on emerging datasets, the NSCI seeks to achieve capable exascale computing by the mid-2020s by driving R&D and leading adoption of these systems.

E.O. 13702 identifies DOE, with its agency missions for science and nuclear stockpile stewardship, as the lead agency for achieving capable exascale computing. DOE's historical computing milestones include deploying the first terascale and petascale supercomputers, and DOE components have assessed the challenges of exascale computing over the last decade.^{5,6} DOE will develop and deploy exascale-class computing by the mid-2020s that supports its science and national security missions and may also benefit the Nation from an economic and educational standpoint. To ensure that the gains of capable exascale computing are broadly realized, a whole-of-Nation approach has been implemented to identify applications from across Federal agencies that can benefit from exascale class computing.

Near-term DOE activities will focus on the following objectives:

- Collaborate with NSCI agencies to identify a range of applications representing government missions for simulation and a quantitative performance-evaluation methodology for each application;

⁵ ACSAC reports and associated materials are available at <http://science.energy.gov/ascr/ascac/reports/>

⁶ The Report of the Task Force on Next Generation High Performance Computing, the Energy Secretary's response, and DOE's assessment of the report are all available at <http://www.energy.gov/seab/downloads/report-task-force-next-generation-high-performance-computing>

- Collaborate with industry to identify and develop solutions to the technical challenges presented by the applications required to support agency missions in capable exascale computing;
- Acquire and deploy next-generation HPC systems, such as the Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) class systems, to explore the technical challenges presented by O(100) petaflop computing and to analyze and solve the most demanding DOE and NNSA mission problems;
- Provide research leadership in the development and use of next-generation HPC computational methods, algorithms, system software, and sustainable application development for DOE missions; and
- Coordinate with the DOD to ensure that future technologies developed for HPC are incorporated into capable exascale systems as appropriate.

Box 2: Capable Exascale Computing Applications

A new 3-year DOE/NIH pilot project exemplifies the interagency collaborations that will contribute to achieving capable exascale computing. Scientists at DOE National Laboratories and NIH's National Cancer Institute (NCI) will use CORAL class computing to extend cancer science and, ultimately, clinical treatments, by:

- Developing new computational approaches to attack mutant proteins that result in certain types of aggressive cancers. This work may lead to potential cancer therapies.
- Using large-scale computations to accelerate the development of patient-derived laboratory models of cancer. Together, DOE and NCI will explore application of deep learning to combine and extract features for large-scale compound and cancer drug screening.
- Understanding the impact of existing therapies in real-world settings.

In addition to improved targeting of drug therapy for cancer patients, this pilot is designed to push the frontiers of high-performance computing. This project integrates both data and science in new ways, potentially leading to future innovations that shape the architectures for exascale while also transforming cancer treatment.

DOE's efforts will be broadly supported by other agencies and activities within NSCI. In 2015, NIH, DOE and NSF jointly issued a "Request for Information (RFI) on Science Drivers Requiring Capable Exascale High Performance Computing."⁷ Over 100 responses were received from academia and industry, and another 135 responses were provided by the DOE National Laboratories, which together highlight a range of scientific opportunities that impact many government agencies.

DOE-sponsored co-design efforts, where government-application owners and prospective exascale system developers explore the performance implications of proposed hardware and software architectures, are being extended to include NSCI agencies. Co-design⁸ processes fully integrate hardware and software with the intent of maximizing performance of exascale applications.

Other co-design efforts are planned, such as NASA participation in capable exascale

system co-design and evaluation to ensure that the special requirements of NASA's mission and diverse application base influence the design of new HPC systems and software.

The other NSCI lead agencies (DOD and NSF) and the foundational R&D agencies (IARPA and NIST) also have key roles in achieving a whole-of-nation effort for exascale computing. For example, in

⁷ The RFI can be found here: <http://grants.nih.gov/grants/guide/notice-files/NOT-GM-15-122.html>

⁸ Co-design processes concurrently develop hardware and software components of complex computer systems with the intent of optimizing overall performance.

addition to coordinating with DOE on advanced architectural design and hardware development, DOD will lead the exploration of computational methods, algorithms, system software, and sustainable application development that exploit NSCI hardware advances and accelerate the development and transition of advanced defense technologies into superior national security capabilities.

NSF will continue to play a central role in scientific discovery, the broader HPC ecosystem for scientific discovery, and workforce development. Overarching NSF activities that are relevant to the capable exascale objective include:

- Identify NSCI-relevant science and engineering frontiers, and characterize the scientific discovery advances accelerated by NSCI activities across NSF-supported disciplines;
- Advance computational and data applications to advance science and engineering, and relevant software technologies; and
- Advance fundamental research in application and system software technologies for programmability and reuse, and for high scalability and fidelity, including methodologies, algorithms, and language-based techniques for accuracy of results, model validation, and uncertainty quantification towards robust and reliable science within the scope of NSCI.

IARPA will support exascale computing by aligning IARPA research efforts on superconducting, machine learning, and post-Moore's Law technologies to pursue the goal of achieving a 100-fold improvement in performance over current computing systems.

NIST will develop the measurement tools and technologies needed to support the next generation of computing technologies and applications through its world-leading research laboratories and strategic collaboration with industry and academia. Specific activities include:

- Provide key enabling platforms to advance the development and testing of novel device architectures and computing platforms;
- Develop measurement science for the physical and materials aspects of future computing technologies with a focus on alternative-state variables that provide low latency, low energy per operation, improved data and communication bandwidth, and higher clock speed;
- Leverage strengths in device physics, materials design, and measurement tool development to address potential logic, memory, storage, and systems concepts needed in an exascale HPC platform;
- Develop measurements, standards, and guidelines for robustness and security of future generation computing systems and networks; and
- Create and evaluate the measurement technologies necessary to assess the reliability and uncertainty of results produced by next-generation computing systems.

In summary, as the path toward capable exascale systems becomes clearer, whole-of-Nation engagement will be utilized to achieve broad use of capable exascale systems across Federal agencies and within the private sector. Success will require increased interagency collaboration and coordination, and collaborations with industry and academia, as appropriate.

3.2. Strategic Objective 2: Technology Coherence

Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.

Box 3: Convergence of Data Analytic Computing and Modeling and Simulation

Historically, there has been a separation between data analytic computing and modeling and simulation. Data analytics focuses on inferring new information from what is already known to enable action on that information. Modeling and simulation focuses on insights into the interaction of the parts of a system, and the system as a whole, to advance understanding in science and engineering and inform policy and economic decision-making. While these systems have traditionally relied on different hardware and software stacks, many of the current challenges facing the two disciplines are similar. The growth of extremely large-scale data analytics within the modeling and simulation community demands a dynamic interaction between analysis and simulations. As demands for computational intensity increase, the data analytics community faces barriers to scalability along with new demands for interoperability, robustness, and reliability of results.

A coherent platform for modeling, simulation, and data analytics would benefit both disciplines while maximizing returns on R&D investments. The primary challenges lie within and across software layers of the HPC environment. In particular, a more agile and reusable HPC software portfolio that is equally capable in data analytics and modeling and simulation will improve productivity, increase reliability and trustworthiness in computations, and establish more sustainable yet agile software. These improvements, in turn, may provide mutual benefit for the analytics and simulation communities, along with new advances across industry, academia, and government.

The benefits of achieving increased coherence between simulation and data analytic computing in HPC applications are common to the NSCI lead agencies and pervade applications across the foundational and deployment agencies. Early engagement with both industry and academia is a common strategy across NSCI agencies to ensure the relevance and impact of the resultant technologies. Initial activities for achieving technology coherence will focus on seed and pilot activities, as well as agencies identifying gaps and opportunities.

Several NSF leadership activities will support improved technology coherence in the use of computing for data analytics and simulation modeling, such as:

- Advances in a diverse, highly interoperable, collaborative, and data-intensive HPC ecosystem in support of NSF science frontiers and broad engagement of academic communities; and
- Advancement of computational and data applications at science and engineering frontiers, as well as their enabling and software technologies.

As part of DOD's leading role in data analytics, DOD will drive the design and development of advanced high performance data analytics capabilities and its supporting ecosystem of software and data science, and will leverage opportunities to increase coherence between the technology base used for modeling and simulation and that used for data analytic computing.

The deployment agencies also have activities contributing significantly to achieving coherence between data analytic computing and modeling and simulation, including:

- NASA will participate in development and implementation of activities to increase synergy between simulation and data-analytic computing, supporting advances in NASA's Big

- Data+Big Compute applications in Earth and space science, aeronautics research, and space exploration;
- NIH will provide research leadership in developing the necessary computational methods, algorithms, and sustainable software applications that will use NSCI technology and advance biomedical research; and
- NOAA will advance its capability to leverage big data as input to its research, modeling, and prediction mission, and as output when delivering innovative products to NOAA's customers.

3.3. Strategic Objective 3: Computing Beyond Moore's Law

Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post-Moore's Law era").

Box 4: Beyond Moore's Law

While current semiconductor approaches should be feasible at exascale, these approaches will eventually plateau due to inherent physical limitations. Strategic Objective 3 looks for new technologies that will continue advancement of computing to establish the foundation to extend the U.S. global computing advantage into the middle of this century. To achieve this goal, the NSCI will pursue two parallel lines of effort over a 10-20 year period: the R&D of technologies that will move digital computing performance past the theoretical limits of complementary metal oxide semiconductors (CMOS) and the research and development of alternative paradigms that will open up new possibilities for the advancement of large-scale computing.

The NSCI envisions a more heterogeneous future computing environment, where digital (von Neumann based) computing is augmented by systems implementing alternative computing paradigms to efficiently solve specific classes of problems. These alternative computational paradigms—whether quantum, neuromorphic or other alternatives—may solve some classes of problems that are intractable with digital computing, and provide more efficient solutions for some other classes of problems. To achieve this vision, it is as essential to research the range of computational problems served by these approaches as it is to further the basic R&D that will enable them. For those problems that will remain in the digital computing domain, increasing performance beyond exascale or further reductions in energy cost may not be feasible with current semiconductor technologies. To address these challenges, NSCI activities must explore and create new highly scalable, programmable, power efficient, and economically viable computing technologies.

The NSCI strategy for future computing recognizes that it is still too soon to pick individual technology paths for alternative computing paradigms or alternative digital computing technologies. The focus of the NSCI over the next 3-5 years will instead be on broad research to identify areas of promise across the span of future computing technologies. For example, superconducting computing is one leading candidate for advanced digital computing technologies, but other yet-unimagined low-energy-per-operation, low latency, and high-speed computing platforms may be possible. The advances required to implement alternative paradigms with a particular technology may be even more daunting. As certain technologies begin to emerge that show both the promise and a reasonable path to groundbreaking capabilities, investment strategies will shift into those areas that will lead to the development of those break-out capabilities.

DOE, IARPA, NIST, NSF, and DOD will all support the development of non-CMOS technologies and non-classical CMOS technologies, with early efforts by NSF, NIST, and IARPA.

The NSCI lead and foundational R&D agencies all play significant roles in accomplishing this strategic objective by exploring alternative computing paradigms or the underlying technologies. NIH, DOE, DOD, and NSF are exploring neuromorphic computing. While not limited to neuromorphic technologies, the National Nanotechnology Initiative's first Grand Challenge seeks to achieve brain-like sensing and problem-solving at low power.⁹ NIST and NSF have long supported research in quantum sensors and quantum communication. NIST, NSF, and DOE have long histories of research in quantum physics, while DOD and IARPA have sponsored basic research in quantum computing for several decades. DOE, IARPA, and NASA are also exploring adiabatic quantum systems. These agencies efforts are supplemented by the NSTC's Quantum Information Science Interagency Working Group, which seeks to develop the scientific basis, infrastructure, future technical workforce, and intellectual property needed to meet agency missions and establish an effective quantum information science industrial sector.

NSF will play a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development required for this future computing environment to:

- Ensure the exploration of diverse scientific inquiry and opportunities at the scientifically rich and rapidly advancing nexus that drives the opportunities for future HPC;
- Foster the use of novel devices and cutting edge technologies to address the scientific needs at the frontiers of knowledge; and
- In collaboration with other NSCI participating agencies, develop a portfolio of foundational research programs to reach the NSCI goal of establishing a viable path forward for future HPC systems even in the post-Moore's Law era.

NSF will expand the scope of its support for foundational research across a range of alternative approaches through four directorates: Engineering (ENG), Computer & Information Science & Engineering (CISE), Mathematical & Physical Sciences (MPS), and Biological Sciences (BIO).

NSF will build on existing interagency and industry partnerships to advance future computing priorities. For example, NSF published a solicitation jointly with industry in December 2015 for "Energy-Efficient Computing: from Devices to Architectures (E2CDA)" to address the NSCI and exascale challenges for power consumption.¹⁰ This partnership will specifically support new research to minimize the energy impacts of processing, storing, and moving data within future computing systems. In addition to targeting an important technical area, this solicitation exemplifies a model for public-private collaboration that may be replicated by other agencies and private sector entities to address other technical requirements.

As foundational R&D agencies for the NSCI, IARPA will increase investments in future computing paradigms offering an alternative to standard semiconductor computing technologies and NIST will develop the measurement tools and technologies needed to support all aspects of future computing technologies.

In particular, IARPA will:

⁹ <https://www.whitehouse.gov/blog/2015/10/15/nanotechnology-inspired-grand-challenge-future-computing>

¹⁰ NSF Program Solicitation 16-256, see <http://www.nsf.gov/pubs/2016/nsf16526/nsf16526.htm>.

- Continue to provide leadership in foundational research on an alternative to standard semiconductor computing technologies;¹¹
- Align research efforts on superconducting, quantum, neuromorphic, and machine learning technologies to efficiently implement applications that are impractical or impossible using digital computing paradigms;¹² and
- Invest in other, post-Moore’s Law technologies that have the potential to support the NSCI Strategic Objectives, as they are identified.

Example NIST activities, executed through its research laboratories and strategic collaboration with industry and academia, include:

- Provide key enabling platforms to advance the development and testing of novel device architectures and computing platforms;
- Develop measurement science for the physical and materials aspects of future computing technologies with a focus on alternative state variables that provide low latency, low energy per operation, improved data/communication bandwidth, and higher clock speed;
- Leverage strengths in device physics, materials design, and measurement tool development to address potential logic, memory, storage, and systems concepts needed in an exascale HPC platform;
- Develop the measurement science to support alternative computational paradigms;
- Continually evaluate the technology pathways, tools, and measurement science that are required to support non-traditional computational paradigms and the classes of problems they may efficiently solve; and
- Create and evaluate the measurement technologies necessary to assess the reliability and uncertainty of results produced by future computing systems.

Among the deployment agencies, NASA will participate in the post-Moore’s Law research, engaging NASA expertise in quantum computing, nanotechnology, and other relevant technologies.

3.4. Strategic Objective 4: An Enduring National HPC Ecosystem

Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.

¹¹ For example, IARPA initiated the Cryogenic Computing Complexity (C3) program in 2013 to address the challenges of “insufficient memory, insufficient integration density, and no realization of complete computing systems” with the goal of establishing superconducting computing as a long-term solution to the power- cooling problem and a successor to end-of-roadmap CMOS for HPC. See <http://www.iarpa.gov/index.php/research-programs/c3>

¹² That is, perform digital computing more efficiently or with higher performance than possible using the capable exascale systems developed under Strategic Objective 1.

Box 5: Envisioning an HPC Ecosystem

Re-envisioning the HPC ecosystem is an integral and enabling element of a broader, richer, and dynamic scientific enterprise and presents a profound opportunity and challenge. With the advent of new instruments and sensors, data science, modeling and simulation, high-speed networking, HPC, and new models of discovery, a dramatic transformation of the scientific process is underway. Change is already apparent within research domains and communities that are mature in their use of HPC, as well as important national programs such as the BRAIN, Materials Genome, and Advanced Manufacturing Initiatives. This transformation will accelerate and impact new communities as data science matures.

For HPC to be effective within such a dynamic, data-rich, and collaborative ecosystem, it must be easily accessible to and usable by the broadest range of researchers. This will require increased investments in people, skills training and outreach to new users, applications, and a variety of innovative technologies as well as creative approaches to ensure that overall productivity of HPC ecosystem-enabled science and engineering research is maximized and sustained.

Greater accessibility and interweaving with production environments will also increase the exposure of HPC capabilities to the full spectrum of cyber threats. Pervasive cybersecurity capabilities that are fully integrated across the HPC ecosystem are needed, but must be incorporated consistently with core American principles of privacy and civil liberty. Integrating security and privacy as foundational elements throughout the full lifespan of this initiative is critical to help avoid costly architectural pitfalls when such elements are brought in after the fact, such as those experienced in the growth of the Internet.

Among Federal agencies, NSF will play the leading role in defining and supporting the development of an enduring national HPC ecosystem, and many agencies will contribute to this shared priority. Other NSCI agencies will complement NSF's HPC workforce development activities through mission-specific efforts. Attendees of the White House Workshop on NSCI noted the importance of developing a capable, multidisciplinary workforce, where domain experts have the skills required to employ HPC tools effectively.¹³ NSF, NIST, and DOE share leadership for algorithm development, and these activities will be supplemented by research sponsored by deployment agencies for specific mission needs. NSF, NIST, and IARPA will lead networking R&D with additional coordination by the NITRD program to increase access to capable exascale systems. R&D of workflow processes will be broadly supported at NSF, but tailoring workflow to meet specific mission needs will be the domain of the relevant NSCI agencies and private industry. Ensuring that the innovations in hardware and software are widely available (rather than benefiting only the users of the most powerful systems) is a shared priority for deployment agencies, industry, and academia.

NSF will play a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development. NSF will:

- Provide leadership in learning and workforce development¹⁴ to encompass support of basic HPC training for a broad user community as well as support for career path development for computational and data scientists;
- Increase engagement with industry and academia through existing programs;

¹³ <https://www.ornl.gov/NSCIWorkshop2015/agenda.htm>

¹⁴ Consistent with "[Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020](#)," Interim Report

- Support the broad deployment of NSCI technologies to increase the capacity and capability of the HPC ecosystem, enabling fundamental understanding across frontiers consistent with NSF scientific and engineering priorities; and
- Lead the development of domestic and international collaborations that will advance transformative computational science and engineering with an integrated approach to high-end computing, data, networking, facilities, software, and multidisciplinary expertise, consistent with NSCI strategic objectives.

DOE will endeavor to maximize the use of its exascale class computing systems.

NSCI-enabled computational environments will allow the DOD science and technology and acquisition workforce to solve its most demanding science and technology problems, and address mission needs.

Other deployment agency efforts include:

- NASA will participate in multi-agency efforts to plan and coordinate optimization of its national HPC infrastructure, engaging NASA's experience in productivity-driven HPC, collaborative computing, large-scale data analysis and visualization environments, massive observational data facilities, and robust nationwide networking;
- NIH will participate in multi-agency efforts to plan and coordinate the development of new and evolving NSCI technologies and algorithms and provide research leadership in developing the necessary computational methods, algorithms, and sustainable software applications that will exploit NSCI technology advances and that are in support of the biomedical research community; and
- NOAA will use its experience with advanced HPC technologies garnered from collaborations with DOE and NSF to upgrade its R&D and Operational HPC systems and continue to invest in software engineering for these architectures in order to maximize the performance and portability of its numerical models, which form the cornerstone of its missions in operational weather forecasting, weather and climate research, understanding and predicting space weather, and understanding ecosystems and coastal issues.

Beyond specific agency plans outlined previously, future areas of focus for NSCI agencies may include research, development, and deployment of:

- Innovative, verifiable, scalable, and reusable application software components, architectures, and platforms to both meet mission needs and easily transition to new and evolving NSCI technologies and algorithms;
- A portfolio of new approaches to dramatically increase productivity in the development and use of parallel HPC applications, introduce the use of HPC in dynamic and control workflows, and dramatically increase accessibility for government, industry, and academia;
- Integrated security capabilities that address the unique aspects of HPC and enable the secure interweaving of HPC and standard and production computing environments;
- Privacy and management controls that address the velocity and data aggregation challenges presented by emerging datasets and converged HPC technologies;
- Robust, secure, and high-performing network fabrics; and

- Next-generation comprehensive cybersecurity and cyber resiliency solutions, including those that take advantage of data analytics and machine learning, such as high-accuracy behavioral and heuristic detection, mitigation of malicious cyber and insider threats, self-healing systems, and predictive resource allocation and management.

NSCI agencies may also develop:

- An ambitious workforce development plan to educate the current generation and train the next generation of scientists and engineers to adopt HPC as an effective approach to solving problems of societal importance; and
- A deployment strategy that recognizes and promotes the benefits of a diverse and changing landscape of technology innovation.

3.5. Strategic Objective 5: Advancement through Public-Private Collaboration

Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

Box 6: Public-Private Collaboration

Historically, HPC developments in the United States derived from close collaboration between Federal agencies, industry, and academia. Urgent national security imperatives in cryptography and nuclear weapons motivated early collaborations, such as the late 1940s-early 1950s efforts with the Atomic Energy Commission. Collaboration continued through the era of the vector-based supercomputer to the foundations of modern HPC, driving transformations that presently define the international HPC landscape. Examples include the Message Passing Interface (MPI) and open source tools for visualization and debugging codes, which have become industry and community standards. These collaborations also led to the adoption of supercomputing as an instrument for U.S. innovation.

Public-private collaboration to advance new R&D approaches to HPC has traditionally focused on HPC hardware. For example, the Nanoelectronics Research Initiative (NRI) is a multiyear collaboration with industry to perform basic research to explore paths beyond the standard semiconductor transistor technology. A joint solicitation on energy-efficient computing (E2CDA) noted in section 3.3 provides a more recent example. While most collaborations have focused on hardware, in order to develop future computing technologies, future collaborations will need to expand to include data analytics, tools and techniques for co-design between software and hardware, improved software-optimization tools, and other areas.

To date, public-private collaboration has been largely driven by specific agency needs and focused on particular targets of opportunity. The NSCI will explore ways to optimize collaboration and continue moving the NSCI *whole-of-government* approach toward a *whole-of-Nation* approach that fully engages the creativity of, and maximizes the benefits to the taxpayer and the U.S. government. NSCI-related opportunities for collaboration will primarily be organized by the NSCI agencies and will be highlighted on the NITRD website.

Public-private collaboration will ensure that advances developed as part of the NSCI provide benefit to industry, academia, and U.S. citizens. Citizens will benefit from advancements in medicine, more accurate prediction of weather and other natural phenomena, and the innovative devices enabled by new materials and computing technologies, among others.

Developing public-private collaboration to ensure broad deployment of NSCI-developed capabilities is a common thread among the NSCI objectives. Architectures and software systems that increase coherence between computational and data intensive workflows will facilitate the convergence of

modeling and data analytics. This, in turn, may provide more capabilities to enable the business and scientific enterprises of the future. Breaking through the limitations of Moore's Law is imperative to producing compact and power-efficient systems, bringing current-day HPC capabilities to new sectors, and establishing new frontiers in computing and analytics.

To fully deploy HPC technologies to existing and new communities, it is not enough to research and develop technical solutions for next-generation HPCs. Although there is a long history among the U.S. Government's HPC-intensive agencies of working closely in this area, more remains to be done to ensure U.S. leadership in computing.

In the near term, NSF's Division of Industrial Innovation and Partnerships will continue to stimulate industrial innovation and university-industry partnerships through existing programs such as the Small Business Innovation Research (SBIR) and Small Business Technology Transfer Programs, Industry/University Cooperative Research Centers, Partnerships for Innovation, Grant Opportunities for Academic Liaison with Industry, and Innovation Corps programs.

The Nation will benefit from better access to HPC tools, particularly for small and medium-sized science and technology businesses. Access to HPC tools may help reduce the concept-to-market time for this critical segment of the U.S. economy. To encourage greater access to these tools, some agencies are including NSCI-related topics in their annual SBIR call for proposals, and are providing access to agency HPC systems in support of SBIR contracts.¹⁵

Broader deployment of HPC technology may also accelerate the process of scientific discovery and innovation by providing capabilities to more academic researchers. As a result, more researchers can use computational and data science to advance their fields. NSCI investments in HPC technologies may also lower the cost of entry into HPC for academic institutions, making HPC systems more accessible at the department or group level, with concomitant impacts on university-based research and training.

Agency-specific activities may be supplemented by interaction with and input from the private-sector HPC community, such as:

- Interagency Requirements Surveys: For each strategic objective, as appropriate, the lead agencies will coordinate the identification, integration, and prioritization of NSCI-related capabilities requirements or mission needs across the agencies;
- Data Collection Activities: Lead agencies will coordinate regular and recurring analysis (market surveys) of NSCI-relevant activities by industry and academia in both the U.S. private sector and abroad, as well as net assessments of U.S. Government efforts compared to those of foreign governments;
- NSCI-Oriented Public Workshops: Through coordination with the NSCI EC, NSCI agencies will convene regular technical exchange forums with U.S. industry and academia for the purpose of producing specific goals and outputs that will advance progress along NSCI strategic goals, supplementing public conferences, advisory panels, forums, and other organized gatherings that occur within the United States or international HPC community; and
- Government Requests for Information (RFIs): As necessary, lead agencies will coordinate NSCI-focused requests for information for U.S. Government, industry, and academia as appropriate,

¹⁵ For example, NASA has an NSCI-focused subtopic in its SBIR call, and provides HPC access for SBIR contracts.

consistent with legal guidance. Industry and academia are strongly encouraged to respond to NSCI-related RFIs to ensure that the full range of views are heard.

4. Summary

E.O. 13702 underscores the growing importance of HPC to the United States for scientific discovery and economic competitiveness. To maximize the benefits of computing for scientific discovery and economic competitiveness, HPC technologies must be adopted by new sectors of the economy and new academic disciplines. U.S. leadership in computing has been maintained for more than 60 years through the joint effort of government, industry, and academia. Sustaining that leadership position will be even more challenging in the coming decades, given increasing computing demands, emerging technological challenges, and growing international competition. To be successful, the NSCI must become a whole-of-Nation effort, implemented through a broad array of public-private collaborative activities and building on a foundation of coordinated government programs involving all executive departments and agencies with significant expertise or equities in HPC. Developing new HPC technologies is necessary but insufficient; the United States must also develop a comprehensive technical and scientific approach and accompanying ecosystem to transition HPC research on hardware, system software, development tools, and applications efficiently into development and, ultimately, operations.

This Plan frames a long-term strategy for achieving that vision. During FY 2016, agencies will implement the strategy through previously planned HPC research, development, and deployment activities. Agencies are also preparing for future activities that will support the principles and objectives of the NSCI. The President's Budget for FY 2017 supports NSCI investments across the Federal Government, including key investments within DOE (\$285 million) and NSF (\$33 million).

To fully achieve the goals and vision of the NSCI will require continuing to transition from a whole-of-government to a whole-of-Nation approach, in which government, industry, and academia together seek and exploit opportunities for coordination and collaboration in the development and use of HPC technologies.

A. Annex A: Related Initiatives

NSCI is intended to support the Administration's priorities and spur innovation. This annex explores how NSCI affects different technology sectors and society by looking at some related initiatives that are national, multi-agency, and have a synergistic relationship to NSCI.

A.1 Advanced Manufacturing Initiatives

The President's proposals to revitalize American manufacturing recognize that a growing and vibrant manufacturing sector is central to our ability to innovate, to export, and to create good-paying American jobs. At the heart of the Administration's efforts for advanced manufacturing is the National Network for Manufacturing Innovation, launched in March 2012 with the goal of catalyzing a network of Manufacturing Innovation Institutes around the country. The Manufacturing Innovation Institutes bring together industry, academia, Federal agencies, State and local governments, and other stakeholders to accelerate innovation by investing in relevant manufacturing technologies with broad applications. In 2014, DOD issued a solicitation for an Integrated Photonics Manufacturing Innovation Institute seeking manufacturing innovations in several areas of this technology following a public-private partnership, cooperative agreement model. Advances in integrated photonics technologies will provide new capabilities and greater performance for future computing and communications platforms. Another source of industry involvement relevant to the NSCI is in the DOE's high-performance computing for manufacturing effort (HPC4Mfg). Through HPC4Mfg, National Laboratory experts in advanced modeling/simulation and data analysis will collaborate with industry on project teams to address manufacturing challenges. More information is available at www.manufacturing.gov.

A.2 BRAIN Initiative

Since its launch in April 2013, the President's BRAIN Initiative—Brain Research through Advancing Innovative Neurotechnologies—has grown to include investments from five Federal agencies: the Defense Advanced Research Projects Agency, NIH, NSF, IARPA, and the Food and Drug Administration. Federal agencies are supporting the Initiative by investing in promising research projects aimed at revolutionizing understanding of the human brain, developing novel technologies, and supporting further R&D in neurotechnology. The President's Budget for FY 2017 also proposes funding for the DOE to join in advancing the goals of the BRAIN Initiative. Major foundations and private research institutions have committed nearly \$500 million to the BRAIN Initiative. In addition, companies have joined this effort through commitments of more than \$30 million in R&D investments. High-performance computing systems, together with applied mathematics for algorithm development, may help advance BRAIN Initiative efforts. As the BRAIN Initiative advances our understanding of cognition and computation in the brain, it will lead to insights about the algorithms used by the brain, knowledge that will help to revolutionize machine learning and, perhaps, even lead to entirely new approaches to computing. More information is available at www.whitehouse.gov/BRAIN.

A.3 Computer Science for All

This year, the President unveiled his plan to give all K-12 students across the country the chance to learn computer science (CS) in school. This initiative builds on a growing movement led by parents, teachers, districts, states, and the private sector to expand CS education. More than \$135 million in investments will be made by NSF and the Corporation for National and Community Service CNCS to

support and train CS teachers across the country over the next 5 years. In addition, in his budget, the President has called for \$4 billion in funding for states, and \$100 million directly for districts to train teachers, expand access to high-quality instructional materials, and build effective regional partnerships. Since the Initiative's launch, twelve states have taken action to expand access to CS education, the private sector has made more \$250 million in philanthropic commitments, and more than 25 Governors have called on Congress to increase K-12 CS funding.

A.4 Materials Genome Initiative

The Materials Genome Initiative (MGI) is a multi-agency initiative designed to create a new era of policy, resources, and infrastructure that supports U.S. institutions in the effort to discover, manufacture, and deploy advanced materials twice as fast, at a fraction of the cost. For many years, the United States has been a dominant player in the discovery of transformative materials that are the basis for entirely new products and industries; yet the time lag between discovery of advanced materials and their use in commercial products can be 20 years or more. MGI will significantly benefit from NSCI, and MGI can also be a valuable source for tools and communities to help guide the NSCI path forward. A central idea of MGI is to provide the seamless integration of data, computational methods, and experimental methods to equip the materials community with advanced tools and techniques to work across materials classes and facilitate development from discovery to manufacturing. More information is available at www.mgi.gov.

A.5 National Big Data R&D Initiative

In March 2012, the Administration announced the National Big Data Research and Development Initiative¹⁶ to leverage the fast growing volumes of digital data to help solve some of the Nation's most pressing challenges and accelerate the Nation's ability to draw insights from large and complex collections of digital data. The initial commitment of more than \$200 million was aimed at improving the tools and techniques needed to access, organize, analyze, and make discoveries from huge volumes of digital data. The goals of the initiative include advancing state-of-the-art core technologies needed to collect, store, preserve, manage, analyze, and share huge quantities of data; harnessing these technologies to accelerate the pace of discovery in science and engineering and transform teaching and learning; and expanding the workforce needed to develop and use big data technologies. NITRD released the Federal Big Data Research and Development Strategic Plan in May 2016 that lays out seven strategies that represent key areas of importance for Big Data R&D¹⁷.

A.6 National Nanotechnology Initiative

The National Nanotechnology Initiative (NNI) is a collaboration between over 20 Federal agencies and departments with the shared vision that the ability to understand and control matter at the nanoscale will lead to revolutionary new materials and devices that will benefit society and the Nation. To accomplish this vision, the goals of the NNI support cutting-edge research in nanotechnology; foster the transfer of nanotechnology innovations from the lab to market; develop and sustain world-class characterization and fabrication facilities; train a skilled workforce to

¹⁶ *National Big Data Research and Development*,

https://www.whitehouse.gov/sites/default/files/microsites/ostp/big_data_press_release_final_2.pdf

¹⁷ The full report can be found here:

https://www.whitehouse.gov/sites/default/files/microsites/ostp/NSTC/bigdatardstrategicplan-nitrd_final-051916.pdf

support nanotechnology innovation; and promote the responsible development of nanotechnology. The 11 agencies with budgets supporting the NNI spend approximately \$1.5 billion per year in nanotechnology R&D, with over \$22 billion invested since 2001. A major thrust of the NNI is the development of nanoscale fabrication processes and materials for use in novel electronic devices and computing systems with lower power demands and increased performance that exceed or extend that of conventional CMOS electronics. This effort includes a significant investment in establishing and sustaining unique user facilities for nanofabrication and materials and device characterization, including the NIST Center for Nanoscale Science and Technology, the NSF-sponsored National Nanotechnology Coordinated Infrastructure centers, and DOE's Nanoscale Science Research Centers. The NNI has fostered several public-private collaborations that could support the NSCI, including the Nanoelectronics Research Initiative and the Semiconductor Technology Advanced Research Network.

The first nanotechnology-inspired Grand Challenge, announced in October 2015, is focused on future computing. The challenge is to create a new type of computer that can interpret and learn from data, solve unfamiliar problems using what it has learned, and identify and solve problems without being asked, while operating with the energy efficiency of the human brain. This challenge will require the development of new kinds of nanoscale devices and materials integrated into three-dimensional systems, and will be implemented in direct coordination with the NSCI. More information is available at www.nano.gov.

A.7 Precision Medicine Initiative

In his 2015 State of the Union address, President Obama launched the Precision Medicine Initiative (PMI) as a bold new effort to revolutionize the practice of medicine by taking into account an individual's genes, environment, and lifestyle to create specific and targeted approaches to treating illness and understanding health. Through collaborative public and private efforts, PMI will leverage advances in genomics, emerging methods for managing and analyzing large data sets while protecting privacy, and health information technology to accelerate biomedical discoveries. Because PMI will enable participant volunteers to share core data, including their electronic health records, health survey information, and mobile health data on lifestyle habits and environmental exposures, it will also focus on extreme-scale data management and analytics. This necessitates leveraging an exascale computing environment with robust privacy and security safeguards in place. Downstream economic benefits of PMI will be realized by improving access to (and enhancement of) biomedical data and information, which has the potential to significantly reduce the lag time between basic research discoveries and tangible clinical outcomes. More information is available at www.wh.gov/PMI.

B. Annex B: E.O. 13702

Executive Order -- Creating a National Strategic Computing Initiative

EXECUTIVE ORDER

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

Section 1. Policy. In order to maximize the benefits of HPC for economic competitiveness and scientific discovery, the United States Government must create a coordinated Federal strategy in HPC research, development, and deployment. Investment in HPC has contributed substantially to national economic prosperity and rapidly accelerated scientific discovery. Creating and deploying technology at the leading edge is vital to advancing my Administration's priorities and spurring innovation. Accordingly, this order establishes the National Strategic Computing Initiative (NSCI). The NSCI is a whole-of-government effort designed to create a cohesive, multi-agency strategic vision and Federal investment strategy, executed in collaboration with industry and academia, to maximize the benefits of HPC for the United States.

Over the past six decades, U.S. computing capabilities have been maintained through continuous research and the development and deployment of new computing systems with rapidly increasing performance on applications of major significance to government, industry, and academia. Maximizing the benefits of HPC in the coming decades will require an effective national response to increasing demands for computing power, emerging technological challenges and opportunities, and growing economic dependency on and competition with other nations. This national response will require a cohesive, strategic effort within the Federal Government and a close collaboration between the public and private sectors.

It is the policy of the United States to sustain and enhance its scientific, technological, and economic leadership position in HPC research, development, and deployment through a coordinated Federal strategy guided by four principles:

1. The United States must deploy and apply new HPC technologies broadly for economic competitiveness and scientific discovery.
2. The United States must foster public-private collaboration, relying on the respective strengths of government, industry, and academia to maximize the benefits of HPC.
3. The United States must adopt a whole-of-government approach that draws upon the strengths of and seeks cooperation among all executive departments and agencies with significant expertise or equities in HPC while also collaborating with industry and academia.
4. The United States must develop a comprehensive technical and scientific approach to transition HPC research on hardware, system software, development tools, and applications efficiently into development and, ultimately, operations.

5. This order establishes the NSCI to implement this whole-of-government strategy, in collaboration with industry and academia, for HPC research, development, and deployment.

Sec. 2. Objectives. Executive departments, agencies, and offices (agencies) participating in the NSCI shall pursue five strategic objectives:

1. Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
2. Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
3. Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post- Moore's Law era").
4. Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.
5. Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

Sec. 3. Roles and Responsibilities. To achieve the five strategic objectives, this order identifies lead agencies, foundational research and development agencies, and deployment agencies. Lead agencies are charged with developing and delivering the next generation of integrated HPC capability and will engage in mutually supportive research and development in hardware and software, as well as in developing the workforce to support the objectives of the NSCI. Foundational research and development agencies are charged with fundamental scientific discovery work and associated advances in engineering necessary to support the NSCI objectives. Deployment agencies will develop mission-based HPC requirements to influence the early stages of the design of new HPC systems and will seek viewpoints from the private sector and academia on target HPC requirements. These groups may expand to include other government entities as HPC-related mission needs emerge.

(a) Lead Agencies. There are three lead agencies for the NSCI: the Department of Energy (DOE), the Department of Defense (DOD), and the National Science Foundation (NSF). The DOE Office of Science and DOE National Nuclear Security Administration will execute a joint program focused on advanced simulation through a capable exascale computing program emphasizing sustained performance on relevant applications and analytic computing to support their missions. NSF will play a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development. DOD will focus on data analytic computing to support its mission. The assignment of these responsibilities reflects the historical roles that each of the lead agencies have played in pushing the frontiers of HPC, and will keep the Nation on the forefront of this strategically important field. The lead agencies will also work with the foundational research and development agencies and the deployment agencies to support the objectives of the NSCI and address the wide variety of needs across the Federal Government.

(b) Foundational Research and Development Agencies. There are two foundational research and development agencies for the NSCI: the Intelligence Advanced Research Projects Activity (IARPA)

and the National Institute of Standards and Technology (NIST). IARPA will focus on future computing paradigms offering an alternative to standard semiconductor computing technologies. NIST will focus on measurement science to support future computing technologies. The foundational research and development agencies will coordinate with deployment agencies to enable effective transition of research and development efforts that support the wide variety of requirements across the Federal Government.

(c) Deployment Agencies. There are five deployment agencies for the NSCI: the National Aeronautics and Space Administration, the Federal Bureau of Investigation, the National Institutes of Health, the Department of Homeland Security, and the National Oceanic and Atmospheric Administration. These agencies may participate in the co-design process to integrate the special requirements of their respective missions and influence the early stages of design of new HPC systems, software, and applications. Agencies will also have the opportunity to participate in testing, supporting workforce development activities, and ensuring effective deployment within their mission contexts.

Sec. 4. Executive Council. (a) To ensure accountability for and coordination of research, development, and deployment activities within the NSCI, there is established an NSCI Executive Council to be co-chaired by the Director of the Office of Science and Technology Policy (OSTP) and the Director of the Office of Management and Budget (OMB). The Director of OSTP shall designate members of the Executive Council from within the executive branch. The Executive Council will include representatives from agencies with roles and responsibilities as identified in this order.

(b) The Executive Council shall coordinate and collaborate with the National Science and Technology Council established by Executive Order 12881 of November 23, 1993, and its subordinate entities as appropriate to ensure that HPC efforts across the Federal Government are aligned with the NSCI. The Executive Council shall also consult with representatives from other agencies as it determines necessary. The Executive Council may create additional task forces as needed to ensure accountability and coordination.

(c) The Executive Council shall meet regularly to assess the status of efforts to implement this order. The Executive Council shall meet no less often than twice yearly in the first year after issuance of this order. The Executive Council may revise the meeting frequency as needed thereafter. In the event the Executive Council is unable to reach consensus, the Co-Chairs will be responsible for documenting issues and potential resolutions through a process led by OSTP and OMB.

(d) The Executive Council will encourage agencies to collaborate with the private sector as appropriate. The Executive Council may seek advice from the President's Council of Advisors on Science and Technology through the Assistant to the President for Science and Technology and may interact with other private sector groups consistent with the Federal Advisory Committee Act.

Sec. 5. Implementation. (a) The Executive Council shall, within 90 days of the date of this order, establish an implementation plan to support and align efforts across agencies in support of the NSCI objectives. Annually thereafter for 5 years, the Executive Council shall update the implementation plan as required and document the progress made in implementing the plan, engaging with the private sector, and taking actions to implement this order. After 5 years, updates to the implementation plan may be requested at the discretion of the Co-Chairs.

(b) The Co-Chairs shall prepare a report each year until 5 years from the date of this order on the status of the NSCI for the President. After 5 years, reports may be prepared at the discretion of the Co-Chairs.

Sec. 6. Definitions. For the purposes of this order:

The term "high-performance computing" refers to systems that, through a combination of processing capability and storage capacity, can solve computational problems that are beyond the capability of small- to medium-scale systems.

The term "petaflop" refers to the ability to perform one quadrillion arithmetic operations per second.

The term "exascale computing system" refers to a system operating at one thousand petaflops.

Sec. 7. General Provisions. (a) Nothing in this order shall be construed to impair or otherwise affect:

1. the authority granted by law to an executive department, agency, or the head thereof; or
2. the functions of the Director of OMB relating to budgetary, administrative, or legislative proposals.

(b) This order shall be implemented consistent with applicable law and subject to the availability of appropriations.

(c) This order is not intended to, and does not, create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person.